UNITED STATES PATENT APPLICATION

OF

NATAN VISHLITZKY, HAIM KOPYLOVITZ AND ELI SHAGAM

FOR

DIGITAL DATA STORAGE SUBSYSTEM INCLUDING ARRANGEMENT FOR INCREASING CACHE

MEMORY ADDRESSABILITY

1   **FIELD OF THE INVENTION**

2   The invention relates generally to the field of digital computer systems and more particularly

3   to a digital data storage subsystem including an arrangement for increasing cache memory

4   addressability.

5   **BACKGROUND OF THE INVENTION**

6   In modern "enterprise" computing environments, that is, computer systems for use in an

7   office environment in a company, a number of personal computers, workstations, mini-computers

8   and mainframe computers, along with other devices such as large mass storage subsystems, network

9   printers and interfaces to the public telephony system, may be interconnected to provide an

10   integrated environment in which information may be shared among users in the company. Typically,

11   users may be performing a variety of operations, including order receipt, manufacturing, shipping,

12   billing, inventory control, and other operations, in which sharing of data on a real-time basis may

13   provide a significant advantage over, for example, maintaining separate records and attempting to

14   later reconcile them. The users may operate on their own data, which they may maintain on the

15   computers they are using, or alternatively they may share data through the large mass storage

16   subsystems.

17   One such large mass storage subsystem is described in, for example, U. S. Patent No.

18   5,206,939, entitled System And Method For Disk Mapping And Data Retrieval, issued April 27,

19   1993 to Moshe Yanai, et al (hereinafter, "the '939 patent"), and U. S. Patent No. 5,381,539, entitled

20   "System And Method For Dynamically Controlling Cache Management," issued January 10, 1995,

21   to Moshe Yanai, et al., both of which are assigned to the assignee of the present invention and

22   incorporated herein by reference. That patent and those applications generally describe an

23   arrangement which allows data, as used by computers, organized in records, with each record being

24   in well-known "CKD" ("count-key-data") format, to be stored in storage devices which provide a

1   "fixed block" storage architecture. In this arrangement, a large cache is used to buffer data that is

2   transferred from the storage devices for use by the respective computers, and, if the data has been

3   modified, transferred back from to the storage devices for storage.

4   In the systems described in the aforementioned patent and patent applications, a directory

5   table is used to provide information concerning the data that is stored in the mass storage subsystem.

6   In one embodiment, in which the mass storage subsystem stores data on a number disk storage

7   devices, the table includes information concerning selected characteristics of each of the CKD

8   records stored in the mass storage subsystem, organized by device, cylinder and read/write head or

9   track, and includes such information as record size and certain formatting characteristics.

10   As disk storage devices increase in storage capacity, it is desirable to increase the size of the

11   cache, the size of the memory that is available for the directory table, and so forth. If these elements

12   are not increased, the cache may become a bottleneck to performance of the mass storage subsystem.

13   A problem arises, however, in that typically, the word size that is used by many processing devices

14   and systems is on the order of thirty-two bits, which would limit addressing of the memory used for

15   the cache, directory table and so forth to four gigabytes (2^{32} bytes).

SUMMARY OF THE INVENTION

17   The invention provides a new and improved mass digital data storage subsystem that includes

18   an arrangement for increasing cache memory addressability.

19   In brief summary, the invention provides a memory manager for use in connection with a

20   memory, the memory manager comprising a memory access request receiver module, an address

21   translation module and a memory access operation control module. The memory access request

22   receiver module is configured to receive an access request requesting an access operation in

23   connection with the memory, the access request including an address. The address translation

24   module is configured to,

1 (i) if the access request requests an access operation in connection with one of a
2   plurality of sections of the memory, provide the address in the access request as an
3   absolute address identifying at least one storage location in the one section; and

4 (ii) if the access request requests an access operation in connection with another of said
5   sections, generate an absolute address from the address provided in the access
6   request, the address in the access request including a segment identifier and an offset,
7   the address translation module being configured to process the segment identifier and
8   offset to generate an absolute address identifying at least one storage location in the
9   other of said sections.

10 The memory access operation control module is configured to perform an access operation in
11 connection with the memory using the absolute address generated by the address translation module.

12 In one embodiment, the memory manager is used in connection with a large memory that is
13 divided into a plurality of sections. One of the sections is a common memory that can be accessed
14 by a number of elements using an absolute address. Another of the sections is in the form of a cache
15 memory divided into a plurality of cache slots, with each cache slot having a plurality of storage
16 locations. In an access request that requests access to a storage location in a cache slot, the segment
17 identifier identifies the cache slot and the offset identifies the storage location in the cache slot that
18 is to be accessed. In that embodiment, the memory manager generates an absolute address by shifting
19 the bits comprising the cache slot identifier to higher-order bit positions and adding the offset
20 thereto. This allows for addressing of larger cache memories.

21   **BRIEF DESCRIPTION OF THE DRAWINGS**

22 This invention is pointed out with particularity in the appended claims. The above and
23 further advantages of this invention may be better understood by referring to the following
24 description taken in conjunction with the accompanying drawings, in which:

1   FIG. 1 is a functional block diagram of a digital computer system, including a mass storage

2   subsystem that includes an arrangement for increasing cache memory addressability, constructed in

3   accordance with the invention;

4   FIGS. 2 is a flow chart depicting operations performed by the mass storage subsystem in

5   connection with the invention.

6   **DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT**

7   The invention will be described in connection with a digital computer system 10 depicted

8   in functional block diagram form in FIG. 1. With reference to FIG. 1, computer system 10 includes

9   a plurality of host computers 11(1) through 11(N) (generally identified by reference numeral 11(n))

10   and a digital data storage subsystem 12 interconnected by a common bus 13. Each host computer

11   11(n) may comprise, for example, a personal computer, workstation, or the like which may be used

12   by a single operator, or a multi-user computer system which may be used by a number of operators.

13   Each host computer 11(n) is connected to an associated channel adapter 24(n), which, in turn, is

14   connected to bus 13. Each host computer 11(n) may control its associated channel adapter 24(n) to

15   perform a retrieval operation, in which the channel adapter 24(n) initiates retrieval of computer

16   programs and digital data (generally, "information") from the digital data storage subsystem 12 for

17   use by the host computer 11(n) in its processing operations. In addition, the host computer 11(n)

18   may control its associated channel adapter 24(n) to perform a storage operation in which the channel

19   adapter 24(n) initiates storage of processed data in the digital data storage subsystem 12. Generally,

20   retrieval operations and storage operations in connection with the digital data storage subsystem 12

21   will collectively be referred to as "access operations."

22   In connection with both retrieval and storage operations, the channel adapter 24(n) will

23   transfer access operation command information, together with processed data to be stored during

24   a storage operation, over the bus 13. Access to the bus 13 is controlled by bus access control

25   circuitry which, in one embodiment, is integrated in the respective channel adapters 15(n). The bus

1   access control circuitry arbitrates among devices connected to the bus 13 which require access to the

2   bus 13. In controlling access to the bus 13, the bus access control circuitry may use any of a number

3   of known bus access arbitration techniques.

4   The digital data storage subsystem 12 in one embodiment is generally similar to the digital

5   data storage subsystem described in U. S. Patent No. 5,206,939, entitled System And Method For

6   Disk Mapping And Data Retrieval, issued April 27, 1993 to Moshe Yanai, et al (hereinafter, "the

7   '939 patent"). As shown in FIG. 1, the digital data storage subsystem 12 includes a plurality of

8   digital data stores 20(1) through 20(M) (generally identified by reference numeral 20(m)), each of

9   which is also connected to bus 13. Each of the data stores 20(m) stores information, including

10   programs and data, which may be accessed by the host computers 11(n) as well as processed data

11   provided to the digital data storage subsystem 12 by the host computers 11(n).

12   Each data store 20(m), in turn, includes a storage controller 21(m) and one or more storage

13   devices generally identified by reference numeral 22. The storage devices 22 may comprise any of

14   the conventional magnetic disk and tape storage devices, as well as optical disk storage devices and

15   CD-ROM devices from which information may be retrieved. Each storage controller 21(m) connects

16   to bus 13 and controls the storage of information which it receives thereover in the storage devices

17   connected thereto. In addition, each storage controller 21(m) controls the retrieval of information

18   from the storage devices 22 which are connected thereto for transmission over bus 13, and in one

19   embodiment includes bus access control circuitry for controlling access to bus 13.

20   The digital data storage subsystem 12 also includes a common memory subsystem 30 for

21   caching information during an access operation and event status information providing selected

22   status information concerning the status of the host computers 11(n) and the data stores 20(m) at

23   certain points in their operations. The caching of event status information by the common memory

24   subsystem 30 is described in detail in U. S. Patent Appn. Ser. No. 08/532,240, filed September 25,

25   1995, in the name of Eli Shagam, et al., and entitled Digital Computer System Including Common

26   Event Log For Logging Event Information Generated By A Plurality of Devices (Atty. Docket No.

27   95-034) assigned to the assignee of the present invention and incorporated herein by reference. The

1   information cached by the common memory subsystem 30 during an access operation includes data

2   provided by a host computer 11(n) to be stored on a data store 20(m) during a storage operation, as

3   well as data provided by a data store 20(m) to be retrieved by a host computer 11(n) during a

4   retrieval operation. The common memory subsystem 30 effectively operates as a buffer to buffer

5   information transferred between the host computers and the data stores 20(m) during an access

6   operation.

7   The common memory subsystem 30 includes a cache memory 31, a cache index directory

8   32 and a common memory section 33, which are generally described in U. S. Pat. Appn. Ser. No.

9   07/893,509 filed June 4, 1995, in the name of Moshe Yanai, et al., entitled "System And Method For

10 Dynamically Controlling Cache Management," assigned to the assignee of the present invention and

11 incorporated herein by reference. The cache memory 31 operates as a buffer in connection with

12 storage and retrieval operations, in particular buffering data received from the host computers 11(n)

13 to be transferred to the storage devices for storage, and buffering data received from the data stores

14 20(m) to be transferred to the host computers 11(n) for processing.

15 The cache memory 31 and cache index directory 32 are generally similar to those described

16 in U. S. Patent Application Serial No. 09/378,644, filed August 20, 1999, in the names of Natan

17 Vishlitzky, et al., and entitled "Digital Data Storage Subsystem Including Directory For Efficiently

18 Providing Formatting Information For Stored Records" (hereinafter the "Vishlitzky application"),

19 incorporated by reference, and will not be described in detail herein. Generally, however, the cache

20 memory 31 includes a series of storage locations, which are organized in a series of cache slots (not

21 separately shown). The cache slots in turn, operate as the cache memory's buffer as noted above.

22 The storage locations in each are identified by a series of addresses, with the starting address of a

23 cache slot being identified by a base address. Each cache slot generally has the capacity to store an

24 entire track of data from a storage device 22.

25 The cache index directory 32 operates as an index for the cache slots in the cache memory,

26 and serves to associate 31. The cache index directory 32 includes a plurality of cache index tables

27 (not separately shown) each of which is associated with one of the storage devices 22 in the storage

1 subsystem 12. The structure of the cache index table associated with a storage device 22 is generally

2 similar to that described in the "Vishlitzky application"), incorporated by reference, and will not be

3 described in detail herein. Generally, each cache index table includes a plurality of cylinder

4 descriptors, each of which is associated with one of the cylinders in the storage device 22 associated

5 with the respective cache index table. Each cylinder descriptor, in turn, includes a plurality of track

6 descriptors, each of which is associated with one of the tracks in the cylinder. Each track descriptor

7 includes status information for the associated track, including a cached pointer and a cache slot

8 pointer field. If data retrieved from a track, or to be stored in a track, is in a cache slot in the cache

9 memory 31, the track descriptor associated with the track will include a set cached flag to so indicate

10 and include a pointer to the cache slot that contains the data in the cache slot pointer field.

11 Each of the channel adapters 15(n) and each of the device controllers 21(m) includes a cache

12 manager 16(n) and 23(m), respectively, to access to the cache memory 31, cache index directory 32

13 and common memory section 33. The particular operations performed during an access operation

14 will depend on a number of factors, including the access operation to be performed, whether or not

15 the data from the particular track to be accessed is cached in the cache memory 31, and whether or

16 not the data contained in a cache slot has been modified or updated by a channel adapter's cache

17 manager 16(n) during a storage operation. As described in the aforementioned Shagam application,

18 the channel adapters 15(n) typically perform storage and retrieval operations in connection with data

19 in the cache memory 31, and the device controllers 21(m) perform "staging" and "de-staging"

20 operations to transfer data in the storage devices 22 to the cache memory 31 for buffering (the

21 staging operations) and to transfer data from the cache memory 31 to the storage devices 22 for

22 storage (the de-staging operations). In performing the staging and de-staging operations, the device

23 controllers 21(m) generally transfer data to and from the cache memory 31 in units of a track, that

24 is, they will during a staging operation transfer all of the data in a track from a storage device 22 to

25 a cache slot in the cache memory 31, and during a de-staging operation copy all of the data in a slot

26 in the cache memory 31 to the track of the storage device 22 from which it was originally staged.

27 The common memory section 33 maintains a number of work lists which are used to control

28 operations by the channel adapters 15(n) and storage controllers 21(m) during an access operation.

1   In particular, the common memory section 33 includes a cache slot replacement list, a pending write

2   list and various lists which the channel adapters 15(n) and storage controllers 21(m) use to

3   communicate to coordinate staging operations (not shown). It will be appreciated that the various

4   lists maintained by the common memory section 33 may comprise any of a number of convenient

5   forms, including queues, trees, stacks or the like. The cache slot replacement list is used to control

6   re-use of cache slots during staging operations in accordance with a convenient cache-slot re-use

7   methodology. During a staging operation, the storage controller's cache manager 23(m) uses the

8   cache slot replacement list to select a cache slot into which it will load the data retrieved from a

9   storage device. (The aforementioned Shagam application describes a modified least-recently-used

10   cache-slot re-use methodology used in one embodiment of the invention). The pending write list

11   is used to identify cache slots which contain updated data, which has not been written to a storage

12   device. During de-staging operations, the storage controllers' cache managers 23(m) will use the

13   write pending list to identify cache slots to be written to a storage device 22. Preferably, the cache

14   slots which are identified in the pending write list will not also be listed in the cache slot replacement

15   list, so that cache slots which contain updated data will not be used until the data has not been

16   written to a storage device through a de-staging operation.

17   The staging operation coordination communication lists include a plurality of stage request

18   lists and a plurality of stage completion lists, with one stage request list being associated with each

19   data store 20(m) and one stage completion list being associated with each host computer 11(n). The

20   channel adapters' cache managers 16(m) use the stage request lists to store stage requests to be

21   performed by the respective data stores 20(m), and the data stores' cache managers 23(n) use the

22   stage completion lists to store stage completion messages to indicate to the respective channel

23   adapters' cache managers 16(m) that the stage requests have been completed.

24   Generally, a channel adapter 24(n), during a retrieval operation, attempts to retrieve the data

25   from the cache memory 31. However, if the data is not in the cache memory 31, it will enable the

26   device controller 21(m) which controls the storage device 22 that contains the data to be retrieved

27   to "stage" the track which contains the data to be retrieved, that is, to transfer all of the data in the

28   track which contains the data to be retrieved into a slot in the cache memory 31. After the data to

1   be retrieved is in a slot in the cache memory 31, the channel adapter 24(n) will retrieve the data from

2   the slot. Similarly, during a storage operation, the channel adapter 24(n) will determine whether the

3   particular track into which the data is to be written is in a slot in the cache memory 31 and if so will

4   store the data in the slot. However, if the data is not in the cache memory 31, the channel adapter

5   24(n) will enable the cache manager 23(m) and storage controller 21(m) which controls the storage

6   device 22 that contains the track whose data is to be updated to perform a staging operation in

7   connection with the track, thereby to transfer the data in the track into a slot in the cache memory

8   31. After the data from the track has been copied into the cache memory 31, the channel adapter

9   24(n) will update the data in the track.

10   The storage controller 21(m) generally attempts to perform a staging operation in connection

11   with an empty slot in the cache memory 31. However, if the storage controller 21(m) may find that

12   all of the cache slots in the cache memory 31 are filled, it will in any case select one of the slots to

13   be used with the staging operation. Before transferring the data from the track to the selected cache

14   slot, it will determine whether the data in the slot has been updated by a storage operation, and if so

15   copy the data to the storage device 22 in a de-staging operation, and thereafter perform a staging

16   operation to copy the data from the storage device to the selected cache slot. It will be appreciated

17   that the storage controller 21(m) need only perform a de-staging operation in connection with a

18   cache slot if the data in the cache slot has been updated, since if the data in the cache slot not been

19   updated before the slot is re-used (which may occur if the channel adapter 24(n) has only performed

20   retrieval operations therewith), the data in the cache slot corresponds to the data in the storage device

21   22.

22   More specifically, as described in the aforementioned Shagam application, during a retrieval

23   operation, the cache manager 16(n) of the initiating channel adapter 24(n) will initially access the

24   cache index table 32(d) in the cache index directory 32 associated with the storage device 22 in

25   which the data to be retrieved is stored, in particular accessing the track descriptor of the cylinder

26   descriptor to determine, from the condition of the cached flag, whether the data from the track is

27   cached in a cache slot in the cache memory. If the cached flag indicates that data from the track is

1   cached in a cache slot the cache manager 16(n) uses the cache slot pointer to identify the particular

2   cache slot in which the data is cached and retrieves the required data from the cache slot.

3   On the other hand, if the cache manager 16(n) determines from the cached flag that the data

4   from the track is not cached in a cache slot, it will generate a stage request to enable the storage

5   controller 21(m) for the storage device 22 which maintains the data to be retrieved, load the stage

6   request in the stage request queue for the data store 21(m) and notify the storage controller 21(m)

7   that a stage request had been loaded in the stage request queue. At some point after receiving the

8   notification, the storage controller 21(m) will retrieve the stage request and perform a staging

9   operation in response thereto. In performing the staging operation, the storage controller 21(m) will

10   retrieve the data from the requested track, use the above-described cache slot replacement list to

11   select a cache slot, load the data into cache slot and update the track descriptor in the cache index

12   table associated with the storage device 22 to indicate that the data from the track is in the cache

13   slot, in particular setting the cached flag and loading a pointer to the cache slot in the cache slot

14   pointer.

15   After the storage controller 21(m) has completed the staging operation, it will load a staging

16   completed message in the stage completion list in the common memory section 33 associated with

17   the host computer 11(n) which issued the staging request, and notify the host computer's cache

18   manager 16(n) that a stage completed message has been loaded therein. At some point after

19   receiving the notification, the host computer's cache manager 16(n) can repeat the operations

20   performed in connection with the retrieval request as described above, in particular accessing the

21   cache index table in the cache index directory 32 associated with the storage device 22 in which the

22   data to be retrieved is stored, in particular accessing the track descriptor of the cylinder descriptor

23   to determine, from the condition of the cached flag, whether the data from the track is cached in a

24   cache slot in the cache memory and, if so, use the cache slot pointer to identify the particular cache

25   slot in which the data is cached and retrieve the required data from the cache slot. Since at this point

26   the cached flag should indicate that the data from the track is cached in a cache slot, the channel

27   adapter's cache manager 16(n) should be able to complete the retrieval operation.

1   Similar operations occur during a storage operation, in which data in a particular track is
2   updated, with the additional operation of removing the identification of the cache slot containing
3   data to be updated from the replacement list and loading it into the pending write list. During a
4   storage operation, the cache manager 16(n) of the initiating channel adapter 24(n) will initially
5   access the cache index table in the cache index directory 32 associated with the storage device 22
6   in which the data to be updated is stored, in particular accessing the track descriptor of the cylinder
7   descriptor to determine, from the condition of the cached flag, whether the data from the track is
8   cached in a cache slot in the cache memory. If the cached flag indicates that data from the track
9   is cached in a cache slot, the cache manager 16(n) uses the cache slot pointer to identify the
10   particular cache slot in which the data is cached and loads the update data into the cache slot. In
11   addition, the channel adapter's cache manager 16(n) will remove the identification of the selected
12   cache slot from the replacement list to the pending write list so that the cache slot will not be re-
13   used until a de-staging operation has been performed in connection with the cache slot.

14   On the other hand, if the cache manager 16(n) determines from the cached flag that the data
15   from the track is not cached in a cache slot, it will generate a stage request to enable the storage
16   controller 21(m) for the storage device 22 which maintains the data to be retrieved, load the stage
17   request in the stage request queue for the data store 21(m) and notify the storage controller 21(m)
18   that a stage request had been loaded in the stage request queue. At some point after receiving the
19   notification, the storage controller 21(m) will retrieve the stage request and perform a staging
20   operation in response thereto. In performing the staging operation, the storage controller 21(m) will
21   retrieve the data from the requested track, select a cache slot, load the data into cache slot and
22   update the track descriptor in the cache index table associated with the storage device 22 to indicate
23   that the data from the track is in the cache slot, in particular setting the cached flag and loading a
24   pointer to the cache slot in the cache slot pointer.

25   After the storage controller 21(m) has completed the staging operation, it will load a staging
26   completed message in the stage completion queue in the common memory section 33 associated with
27   the host computer 11(n) which issued the staging request, and notify the cache manager 16(n) that
28   a stage completed message has been loaded therein. At some point after receiving the notification,

-11-

1   the cache manager 16(n) can repeat the operations performed in connection with the retrieval request

2   as described above, in particular accessing the cache index table in the cache index directory 32

3   associated with the storage device 22 in which the data to be retrieved is stored, in particular

4   accessing the track descriptor of the cylinder descriptor to determine, from the condition of the

5   cached flag, whether the data from the track is cached in a cache slot in the cache memory and, if

6   so, use the cache slot pointer to identify the particular cache slot in which the data is cached and

7   retrieve the required data from the cache slog. Since at this point the cached flag should indicate

8   that the data from the tack is cached in a cache slot, the cache manager 16(n) should be able to

9   complete the storage operation as described above.

10   As described above, the data stores' cache managers 23(m) also perform de-staging

11   operations using the pending write list to identify cache slots which contain updated data to be

12   written back to the original storage device 22 and track whose data was cached in the respective

13   cache slots. When a cache slot is de-staged, since at that point the data in the cache slot

14   corresponds to the data on the respective storage device 22, the data store's cache manager 23(m)

15   which performs the de-staging operation will remove the cache slot's identification from the pending

16   write list and return it to the replacement list so that the cache slot can be removed. It will be

17   appreciated, however, that a host computer's cache manager 16(n) may perform a number of retrieval

18   operations and/or storage operations in connection with data in the same cache slot after the data

19   in the track cached in the slot has been staged and before it can be de-staged, and so data in a cache

20   slot can be updated a number of times before it is de-staged. In addition, it will be appreciated that

21   after a cache slot has been de-staged, it may also be updated during a storage operation before the

22   cache slot is re-used during a staging operation. When that occurs however, since, as described

23   above, the host computer's cache manager 16(m) removes the cache slot's identification from the

24   replacement list and placed it on the write pending list as part of the storage operation, the cache slot

25   will be subject to another de-staging operation before it can be re-used. Thus, a particular cache slot

26   may be subject to de-staging a number of times with data cached for the same storage device 22,

27   cylinder and track, without being reused.

1   In addition to the work lists, the common memory section 33 may provide storage for other
2   data, including, for example the aforementioned event log, status information, and other information
3   that may be useful in connection with the operation and maintenance of the storage subsystem 12.

4   The invention provides an arrangement for efficiently addressing the common memory
5   subsystem 30. In one embodiment, the various portions of the common memory subsystem 30,
6   namely, the cache memory 31, the cache index directory 32 and the common memory section 33,
7   reside in a commonly-addressable random-access memory. When a cache manager 23(m), 25(n)
8   accesses the common memory section 33, it will generally do so by providing a "linear" address, that
9   is a single address that generally identifies the storage location continuing the data to be retrieved
10   during a retrieval operation, or in which data is to be stored during a storage operation.

11   On the other hand, when a cache manager 23(m), 25(n) accesses the cache memory 31 or the
12   cache index directory 32, it will generally do so by providing a "segmented" address. That is, when
13   a cache manager 23(m), 25(n) accesses the cache memory 31, it will provide a cache slot pointer and
14   an offset. The cache slot pointer identifies the particular cache slot in the cache memory 31 that
15   includes the storage location from which data is to be retrieved during a retrieval operation, or in
16   which data is to be stored during a storage operation. The offset, in turn, identifies the offset, from,
17   for example, the first storage location in the cache slot, which is identified by the base address for
18   the cache slot, to the particular storage location in which data is to be stored or from which data is
19   to be retrieved.

20   Similarly, when a cache manager 23(m), 25(n) accesses the cache index directory 32, it will
21   provide a device pointer and an offset. The device pointer identifies the particular one of the storage
22   devices 22 whose cache index table is to be accessed, and the offset provides an offset from the
23   beginning of the cache index table to the particular storage location in which data is to be stored
24   during a storage operation, or from which data is to be retrieved during a retrieval operation.

25   The common memory subsystem 30 also includes a memory manager 34 that performs a
26   number of operations in connection with the invention. In particular, the memory manager 34
27   receives storage and retrieval requests from the cache managers 23(m), 25(n). If a request is to store

-13-

1   data in, or retrieve data from, a storage location in the common memory section 33, it will make use

2   of the address provided in the request in performing the storage or retrieval operation. In addition,

3   the memory manager 34 may validate the address to verify that the address is within the region of

4   memory for the common memory section 33.

5   On the other hand, if the request is to store data in, or retrieve data from, a storage location

6   in the cache memory 31, the memory manager 34 will generate from the cache slot pointer and offset

7   a linear address that can be used in addressing the cache memory 31. In addition, the memory

8   manager 34 will validate the cache slot pointer to verify that it points to a valid cache slot, and the

9   offset to verify that it points to a valid storage location within the cache slot. In validating the offset,

10   the memory manager 34 can determine whether the value of the offset is larger than the number of

11   storage locations in the cache slot. Generally, all of the cache slots will contain the same number

12   of storage locations, and the memory manager 34 can compare the value of the offset to a value that

13   corresponds to the number of storage locations in a cache slot. If the value of the offset is greater

14   than the number of storage locations in a cache slot, the memory manager 34 will determine that the

15   offset is invalid. On the other hand, if the value of the offset is less than or equal to the number of

16   storage locations in a cache slot, the memory manager 34 will determine that the offset is valid.

17   Similarly, in validating the cache slot pointer, the memory manager 34 can compare the value

18   represented by the cache slot pointer to the number of cache slots. If the value represented by the

19   cache slot pointer is greater than the number of cache slots in the cache memory 31, the memory

20   manager 34 will determine that the cache slot pointer is invalid. On the other hand, if the value

21   represented by the cache slot pointer is less than or equal to the number of cache slots, the memory

22   manager 34 will determine that the cache slot pointer is valid.

23   On the other hand, if the request is to store data in, or retrieve data from, a storage location

24   in the cache index directory 32, the memory manager 34 will generate from the device pointer and

25   offset a linear address that can be used in addressing the cache memory 31. In addition, the memory

26   manager 34 will validate the device pointer to verify that it points to a valid storage device, and the

27   offset to verify that it points to a valid storage location within the storage device's cache index table.

28   In validating the offset, the memory manager 34 can determine whether the value of the offset is

1   larger than the number of storage locations in the cache index table. Generally, all of the cache index

2   tables will contain the same number of storage locations, and the memory manager 34 can compare

3   the value of the offset to a value that corresponds to the number of storage locations in a cache index

4   table. If the value of the offset is greater than the number of storage locations in a cache index table,

5   the memory manager 34 will determine that the offset is invalid. On the other hand, if the value of

6   the offset is less than or equal to the number of storage locations in a cache index table, the memory

7   manager 34 will determine that the offset is valid. Similarly, in validating the device pointer, the

8   memory manager 34 can compare the value represented by the device pointer to the number of

9   storage devices. If the value represented by the device pointer is greater than the number of storage

10   devices 22 in the data stores 20(m), the memory manager 34 will determine that the device pointer

11   is invalid. On the other hand, if the value represented by the device pointer is less than or equal to

12   the number of storage devices 22, the memory manager 34 will determine that the device pointer is

13   valid.

14   In one embodiment, the memory manager 34 generates a linear address from a cache slot

15   pointer and offset, during a storage or retrieval operation in connection with a cache slot, in the same

16   manner that it generates a linear address from the device pointer and offset, during a storage or

17   retrieval operation in connection with a cache index table. Generally, after validating the respective

18   cache slot pointer or device pointer and the offset, the memory manager 34 generates the linear

19   address by shifting the binary representation of the pointer "to the left", by a predetermined number

20   "n" of bit positions, so that the "i-th" bit of the pointer is shifted to the "i+n-th" bit position. If the

21   number of bits in the respective pointer is "I," and bit positions in the respective pointer are labeled

22   I-1 (which corresponds to the left-most, most significant bit position), through "0" (which

23   corresponds to the right-most, least significant bit position), bits in the "n" left-most bit positions I-1

24   through I-1-n of the pointer will be shifted out of the shifted pointer, and bits having selected values,

25   for example "zero's," will be shifted into the "n" right-most bit positions 0 through n-1. After the

26   pointer has been shifted to the left, the offset will be added to the shifted pointer to generate the

27   linear address.

1   In one embodiment, the number "n" of bit positions that the respective cache slot pointer or

2   device pointer will be shifted is selected to be "seven." Providing a shift of seven allows the number

3   of addressable storage locations that can be addressed to be increased by 2^7, or 128. Accordingly,

4   if a cache manager 23(m), 25(n) were to provide a storage or retrieval request to the memory

5   manager 34 that included a thirty-two bit linear address, which could identify a storage location in

6   a common memory subsystem 30 that comprises a 4GB (four gigabyte) memory, by instead having

7   the cache manager 23(m), 25(n) provide a pointer and offset and having the memory manager 34

8   generate a linear address by shifting the cache slot/device pointer by seven bit positions and adding

9   the offset thereto will allow the identification of a storage location in a common memory subsystem

10   30 that comprises a 512GB memory.

11   Operations performed by the memory manager 34 in connection with a retrieval or storage

12   request from a cache manager 23(m), 25(n) will be described in connection with the flow chart

13   depicted in FIG. 2. With reference to FIG. 2, after the memory manager 34 receives a storage or

14   retrieval request from a cache manager 23(m), 25(n) (step 100), it will initially determine whether

15   the request is to store data in or retrieve data from a storage location in the common memory section

16   33, on the one hand, or the cache memory 31 or cache index directory 32, on the other hand (step

17   101). The memory manager 34 can determine in step 101 that the request is to store data in or

18   retrieve data from a storage location in the common memory section 33 if, for example the request

19   includes a single, for example thirty-two bit address, which it can then determine to be a linear

20   address. On the other hand, the memory manager 34 can determine in step 101 that the request is

21   to store data in or retrieve data from the cache memory 31 or cache index directory 32 if, for

22   example, the request includes two addresses, one of which it can determine to be a cache slot or

23   device pointer and the other an offset.

24   If the memory manager 34 determines in step 101 that the request is to store data in or

25   retrieve data from a storage location in the common memory section 33, it will sequence to a step

26   102 in which it will determine whether the address received with the request is within the range of

27   addresses for the common memory section 33. If the memory manager 34 makes a positive

28   determination in step 102, it will use the address in connection with a storage or retrieval operation

-16-

1   in connection with the storage location in the common memory section 33 identified by the address

2   (step 103). Thereafter, the memory manager 34 can provide a notification to the cache manager

3   23(m), 25(n) from which the request was received in step 100 indicating that the respective storage

4   or retrieval operation was successfully completed (step 104).

5   Returning to step 102, if the memory manager 34 makes a negative determination in that

6   step, it will sequence to step 105, in which it will provide a notification to the cache manager 23(m),

7   25(n) from which the request was received in step 100 indicating that the respective storage or

8   retrieval operation was not successfully completed. Thereafter, the cache manager 23(m), 25(n) may

9   perform predetermined error recovery operations.

10   Returning to step 101, if the memory manager 34 determines in that step that the request is

11   to store data in or retrieve data from a storage location in the cache memory 31 or the cache index

12   directory 32, it will sequence to a step 110 in which it will determine whether the pointer provided

13   as part of the request is a valid pointer (step 110). It will be appreciated that, if the request is to store

14   data in or retrieve data from a storage location in the cache memory 31, the memory manager 34 will

15   determine whether the pointer is valid by determining whether it points to a valid cache slot. On the

16   other hand, if the request is to store data in or retrieve data from the cache index directory 32, the

17   memory manager 34 will determine whether the pointer is valid by determining whether it points

18   to a valid storage device 22.

19   In either case, if the memory manager 34 makes a positive determination in step 110, which

20   will be the case if the pointer provided as part of the request points to a valid cache slot or cache

21   index table, it will sequence to step 111 to determine whether the offset provided as part of the

22   request is a valid offset. It will be appreciated that, if the request is to store data in or retrieve data

23   from a storage location in the cache memory 31, the memory manager 34 will determine whether

24   the offset is valid by determining whether the value of the offset is less than the number of storage

25   locations in a cache slot. On the other hand, if the request is to store data in or retrieve data from the

26   cache slot directory 32, the memory manager 34 will determine whether the pointer is valid by

1 determining whether the value of the offset is less than the number of storage locations in a cache

2 index table in the cache index directory 32.

3   If the memory manager 34 makes a positive determination in step 111, which will be the case

4 if the offset provided as part of the request is a valid offset, it will sequence to a series of steps to

5 generate a linear address using the pointer and offset. In that operation, the memory manager 34 will

6 shift the bits comprising the pointer to the "left" by a predetermined number of bit positions (step

7 112), and add the offset to the shifted pointer, thereby to generate the linear address (step 113). After

8 the memory manager 34 has generated the linear address, it will use the linear address in connection

9 with a storage or retrieval operation in connection with the storage location in the respective cache

10 memory 31 or cache index directory 32 as identified by the address (step 114). Thereafter, the

11 memory manager 34 can provide a notification to the cache manager 23(m), 25(n) from which the

12 request was received in step 100 indicating that the respective storage or retrieval operation was

13 successfully completed (step 115).

14   Returning to steps 110 and 111, if the memory manager 34 makes a negative determination

15 in either step, it will sequence to step 105, in which it will provide a notification to the cache

16 manager 23(m), 25(n) from which the request was received in step 100 indicating that the respective

17 storage or retrieval operation was not successfully completed (step 116). It will be appreciated that

18 the memory manager 34 will make a negative determination in step 110 if it determines that the

19 pointer provided as part of the request does not identify a valid cache slot or storage device 22. After

20 the cache manager 23(m), 25(n) receives the notification, it may perform predetermined error

21 recovery operations.

22   The invention provides a number of advantages. In particular, the invention provides an

23 arrangement that allows for the addressing of a random-access memory, such as the common

24 memory subsystem 30 having increased numbers of storage locations.

25   It will be appreciated that a number of modifications may be made to the digital data storage

26 subsystem as described above in connection with FIGS. 1 and 2. For example, although the storage

27 subsystem 12 has been described as being connected to a plurality of host computers 11(n), it will

1   be appreciated that the storage subsystem may be used only in connection with a single host

2   computer 11(n). In addition, although the storage subsystem 12 has been described as including a

3   plurality of data stores 20(m) and a plurality of storage devices 22, it will be appreciated that the

4   storage subsystem 12 may be provided with a single data store and a single storage device.

5   Furthermore, it will be appreciated that, although the memory manager 34 has been described

6   as directly using the pointer provided with the access request in generating an absolute address, it

7   will be appreciated that the memory manager 34 may instead use the pointer provided with the

8   access request as an index into a table. The table, in turn, has a plurality of entries, each of which

9   is associated with one of the values of the pointers that may be provided with an access request and

10   contains a value that, when shifted by the selected number of bit positions, provides the base address

11   for the cache slot or directory table in the common memory subsystem. This will alleviate having

12   the devices that generate the access requests to be provided with such information.

13   In addition, it will be appreciated that the number "n" of bit positions that the respective

14   pointer is shifted may differ from "seven," and may comprise any convenient value.

15   It will be appreciated that a system in accordance with the invention can be constructed in

16   whole or in part from special purpose hardware or a general purpose computer system, or any

17   combination thereof, any portion of which may be controlled by a suitable program.

18   It will be appreciated that a system in accordance with the invention can be constructed in

19   whole or in part from special purpose hardware or a general purpose computer system, or any

20   combination thereof, any portion of which may be controlled by a suitable program. Any program

21   may in whole or in part comprise part of or be stored on the system in a conventional manner, or it

22   may in whole or in part be provided in to the system over a network or other mechanism for

23   transferring information in a conventional manner. In addition, it will be appreciated that the system

24   may be operated and/or otherwise controlled by means of information provided by an operator using

25   operator input elements (not shown) which may be connected directly to the system or which may

26   transfer the information to the system over a network or other mechanism for transferring

27   information in a conventional manner.

1   The foregoing description has been limited to a specific embodiment of this invention. It will

2   be apparent, however, that various variations and modifications may be made to the invention, with

3   the attainment of some or all of the advantages of the invention. It is the object of the appended

4   claims to cover these and such other variations and modifications as come within the true spirit and

5   scope of the invention.

6   What is claimed as new and desired to be secured by Letters Patent of the United States is: